

Heterogeneous Multi-robot System for Exploration and Strategic Water Sampling

Sandeep Manjanna¹, Alberto Quattrini Li², Ryan N. Smith³, Ioannis Rekleitis² and Gregory Dudek¹

Abstract—Physical sampling of water for off-site analysis is necessary for many applications like monitoring the quality of drinking water in reservoirs, understanding marine ecosystems, and measuring contamination levels in fresh-water systems. In this paper, the focus is on algorithms for efficient measurement and sampling using a multi-robot, data-driven, water-sampling behavior, where autonomous surface vehicles plan and execute water sampling using the chlorophyll density as a cue for plankton-rich water samples. We use two Autonomous Surface Vehicles (ASVs), one equipped with a water quality sensor and the other equipped with a water-sampling apparatus. The ASV with the sensor acts as an explorer, measuring and building a spatial map of chlorophyll density in the given region of interest. The ASV equipped with the water sampling apparatus makes decisions in real time on where to sample the water based on the suggestions made by the explorer robot.

We evaluate the system in the context of measuring chlorophyll distributions. We do this both in simulation based on real geophysical data from MODIS measurements, and on real robots in a water reservoir. We demonstrate the effectiveness of the proposed approach in several ways including in terms of mean error in the interpolated data as a function of distance traveled.

I. INTRODUCTION

In this paper we propose and evaluate the design of a multi-robot system composed of two heterogeneous robots – one equipped with a water quality sensor to measure a phenomenon, the other one equipped with a water sampling apparatus for collection of water samples in real time.

Collection of water samples is an essential element of marine science, marine biology, limnology, public health, and related disciplines. While some measurements can be made *in situ* and in real time, many important measurements can only be accomplished by collecting physical samples in the domain of interest and doing the analysis at a suitable remote facility (i.e., “back in the lab”). In many cases, the selection of suitable sampling locations can have a large impact on the quality and accuracy of the estimation process; for example if pollutant extrema are being estimated. Traditional methods for sampling depend heavily on manual labor, are time consuming, and can be fraught with risks of human error. Robotic sampling systems allow scientists to collect




Fig. 1. Two Clearpath Heron ASVs (a), one equipped with a water quality sensor (b), another with a water sampling apparatus (c).

richer and more complete data sets that would normally be impossible using traditional manual data collection [1].

In this paper, we address the problem of estimating a spatially-varying phenomenon over a region and collecting water samples with emphasis on good sampling locations without any prior knowledge of the spatial field. This task is completed by a heterogeneous robotic team, composed of two robotic boats, an *explorer* that measures variables to suggest sample utility and a *sampler* that collects physical samples (Figure 1). Das *et al.* [2] proposed a probabilistic method for a single Autonomous Underwater Vehicle (AUV) that can monitor and sample. In our case, we divide the task between two robots. This provides an efficient trade off between system complexity, payload capacity, and run time, besides improving the quality of the collected samples – where quality is expressed as the sum of measured values of the samples collected.

In particular, such a task leads to two related subproblems: *exploration* and *sampling*. We propose an exploration strategy for the *explorer* that makes real time observations to create a preliminary map and suggest potential locations that are good for sampling. The *sampler* then uses these suggestions to decide on the best location and collects water samples at these locations. Our exploration technique is based on the concept of *frontier*-based exploration, similar to that introduced by Yamauchi [3] for indoor map building and exploration. In this approach, the robot makes exploratory decisions based on the latest information gathered. This frontier-based technique scales well with the size of the region, differently from common coverage approaches that employ a boustrophedonic coverage path [4], [5]. Notice that the absence of prior information on the spatial distribution of the data prevents us from using alternative powerful selective coverage methods [6]–[10].

We design a strategy for the *sampler* to choose among the locations where measurements were taken by the *explorer*,

¹S. Manjanna and G. Dudek are with School of Computer Science, McGill University, Montreal, Quebec, Canada msandeep, dudek@cim.mcgill.ca

²A. Quattrini Li and I. Rekleitis are with the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC 29208, USA albertoq, yiannisr@cse.sc.edu

³R. N. Smith is with Fort Lewis College, Durango, CO 81301, USA. He was supported by NSF MRI 1531322 and Office of Naval Research Award N000141612634. rnsmith@fortlewis.edu

so that the physical sample is also associated with a prior sensor measurement. Note that in the scenario considered, it is not ideal to wait until the end of exploration: first, the physical sample should be collected temporally close to the measurement; second, in this way the total time spent on sampling is optimized. There are many variants and solutions for this problem in the field of spatial sampling. Girdhar *et al.* propose a multi-choice hiring algorithm [11] for making irrevocable hiring decisions from a stream of candidates. Another approach is to use multiple time windows and treat sampling within each window as a classic secretary hiring problem as proposed by Bateni *et al.* in submodular secretary algorithm [12]. We will further discuss the submodular analysis of the secretary algorithm in Section III-B.

There is a body of existing work on using multi-robot systems to explore and map a spatial phenomenon. In [13], Valada *et al.* developed a low-cost multirobot autonomous platform, and tested the proposed system for monitoring water quality. The paper proposes a discretization of the area and a strategy based on maximum uncertainty. Girdhar *et al.* demonstrated a heterogeneous multirobot system composed of Unmanned Aerial Vehicle (UAV), Autonomous Surface Vehicle (ASV), and an AUV covering an area of interest [14], where the regions to cover are given by a human expert. Many systems have been proposed that are capable of collecting water samples. A catamaran with a water sampling system was proposed by Caccia *et al.* [15] and tested near Antarctica. Ore *et al.* [16] presented a UAV equipped with a water sampling apparatus. Robotic physical sampling has also been approached in domains other than marine robotics, such as planetary robots or mining robots to collect samples of rocks, ores, and other terrestrial samples (e.g., [17]). These papers focus more on the hardware design of the sampling platform and the autonomy that allows robots to navigate environments and collect samples, and not on the sampling location selection strategy as proposed in this paper.

Exploration techniques have a key application in addressing search and rescue problems [18], and gas leakage detection [19]. Some methods [6], [20] assume to have a priori information available so that areas can be selectively explored to increase the reward over time. However, in our scenario, prior information is not available and is estimated in real time by the explorer. Our focus is on building a representation of the spatial field and making informed decisions about when and where to collect a water sample.

The paper is organized as follows. An overview of the problem and the proposed methods for the robotic team are presented in the next two sections. In Section IV, we validate the proposed method both in simulations and field experiments. Finally, Section V concludes the paper discussing some of the lesson learned and outlining future work.

II. PROBLEM STATEMENT

Two ASVs are deployed in a continuous two-dimensional area of interest $\mathcal{E} \subset \mathbb{R}^2$ with user-defined boundaries. We assume that such an area is obstacle-free as in many

marine science expeditions. Both the vehicles navigate using differential drive and are equipped with GPS to localize and WiFi channel for communication. The *explorer* ASV is equipped with a water quality sensor and is assigned with the task of exploring the region by building a representation of the spatial phenomenon, and thus suggesting interesting locations where to collect water samples. The *sampler* ASV has a water sampling apparatus with k sampling units to collect water samples to be analyzed a posteriori in the lab. As the mission evolves, the explorer selects a series of destination poses to get more measurements and builds a reliable model of the area by taking measurements at locations with high uncertainty. At the same time, the sampler receives measurements from the explorer and uses this information to decide where to take a sample. The mission progresses up to the mission duration T_m , which generally depends on the specific application. All k units of the water sampling apparatus should be used within the mission duration T_m . Even if the ultimate objective of the multirobot team is to maximize the total value of the collected samples, this process leads to two related problems addressed in this paper:

- 1) *Exploration*: explorer selects a sequence of poses $Q = \langle q_0, q_1, \dots, q_n \rangle$, with $q_i \in \mathcal{E}$, so that the model of the area converges to the true phenomenon. Note that this process is run on-line, and the explorer makes decisions as new measurements y_i associated with GPS locations \mathbf{x}^i are collected. The goal is to improve the quality of the map and reduce the traveled distance.
- 2) *Sampling*: based on all the measurements \mathbf{Y} , the sampler selects a number of locations \mathcal{L} to collect physical samples, where $|\mathcal{L}| = k$ and $l \in \mathcal{L} \iff \exists y^i \in \mathbf{Y} \mid \mathbf{x}^i = l$. The final objective is to maximize the sum of the values at sampled locations ($\sum_{l \in \mathcal{L}^*} f(l)$) within the maximum duration of the mission T_m .

Intuitively, the performance of the sampler can be improved by improving the performance of the explorer.

III. INFORMED STRATEGIC SAMPLING

The proposed system is based on using two robots that coordinate with each other to achieve the ultimate goal of sampling. Frontier-based exploration is used by the explorer, while a variant of the secretary hiring problem is used for the sampler. The following subsections report the details of both components.

A. Gaussian Process Frontier-based Exploration

The explorer's objective is to select locations $L^* = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^m]$ over time such that the phenomenon is mapped efficiently. Note that, while the robot is traveling to those locations, measurements $\mathbf{Y} = [y^1, y^2, \dots, y^t]$ with associated GPS locations $\mathbf{X} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^t]$ are collected. The goal is to optimize the time and the traveled distance, yet create a good model $\hat{f}(\mathbf{x})$ of the spatial phenomenon $f(\mathbf{x})$.

With finite time and finite battery life of the robot, it is not feasible to take measurements at every location in the region of interest \mathcal{E} . Hence, we use Gaus-




Fig. 2. Candidate locations for the explorer generated by two techniques at a mission time step. The colormap represents the variance in the spatial representation of the field. Red circles represent the potential candidate locations l . Black lines show the contours. (a) Contour-based location selection. (b) Fixed-window location selection.

sian Processes (GP) [21] to model the spatial field. In particular, a phenomenon over locations \mathbf{W} can be estimated as a posterior distribution $p(f(\mathbf{W}) \mid \mathbf{W}, \mathbf{X}, \mathbf{Y}) \sim \mathcal{N}(\mu_{\mathbf{W}}, \Sigma_{\mathbf{W}})$ fitted over a set of noisy observations \mathbf{Y} made at locations \mathbf{X} . The mean vector $\mu_{\mathbf{W}}$ is obtained as $\mu_{\mathbf{W}} = \mu(\mathbf{W}) + K(\mathbf{X}, \mathbf{W})^T K(\mathbf{X}, \mathbf{X})^{-1}(\mathbf{Y} - \mu(\mathbf{X}))$ and represents the estimate of the phenomenon, while the covariance matrix is given by $\Sigma_{\mathbf{W}} = K(\mathbf{W}, \mathbf{W}) - K(\mathbf{X}, \mathbf{W})^T K(\mathbf{X}, \mathbf{X})^{-1} K(\mathbf{X}, \mathbf{W})$. Mean and covariance functions should be formulated to completely define a GP. As done in mainstream approach, mean is assumed to be zero, and the covariance function $k(\mathbf{x}, \mathbf{x}')$, is a radial basis kernel (RBF):

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right), \quad (1)$$

where signal variance σ_f^2 and length scale l^2 are hyperparameters that encode amplitude and smoothness. Note that, with a GP, it is possible to quantify the uncertainty of the estimates in \mathbf{W} by looking at the main diagonal of $\Sigma_{\mathbf{W}}$, also called *predictive variance*.

Our exploration technique uses a one-step look ahead, where the robot decides on a set of locations to visit at epoch m only after reaching the chosen location of epoch $m - 1$. We propose two methods to generate a list of locations (Figure 2). One of the approaches is to consider locations on the outer-most contour between a region with high variance and a region with low variance (Figure 2(a)). An easier method is to consider all the locations on a fixed planning window centered on the current position of the robot (Figure 2(b)).

The list of new locations is added to the list of candidate locations L , thus the algorithm chooses among all the locations around the current trajectory of the explorer. The rationale is that every measurement decreases the variance within a window; as such, the ASV should go to the boundary of that window to build an efficient representation of the spatial field with minimum distance traveled. Candidate locations L are then evaluated based on the predicted variance at these locations according to the learned GP model and their distance from the current robot location. The location with highest predicted variance and least distance is chosen as the current step target. We use a normalized blending function (Eq. (2)) to resolve the trade-off between distance




Fig. 3. RMSE of the generated map of Chlorophyll density plotted against distance traveled by the robot. The comparison is between multiple sized fixed windows. Error bars indicate standard deviation over five real time simulation trials. Note that the x-axis starts at 100 m.




Fig. 4. RMSE of the generated map of Chlorophyll density plotted against distance traveled by the robot. The comparison is between fixed window technique and contour based technique. Error bars indicate standard deviation over five real time simulation trials. Note that the x-axis starts at 100 m.

and variance.

$$l^* = \arg \min_{l \in L} ((1 - w(t)) * \tilde{d}(\mathbf{x}_m, l) + w(t) * \tilde{v}(l)), \quad (2)$$

where $\tilde{d}(\mathbf{x}_m, l) = d(\mathbf{x}_m, l) / \max_{l \in L} d(\mathbf{x}_m, l)$ is a normalized distance between the current robot position and candidate location l ; and $\tilde{v}(l) = 1 - (v(l) / \max_{l \in L} v(l))$ is the normalized variance at location l . With time, it is beneficial to explore the locations with higher variance even if they are far from the robot's current location. We need to weigh the variance higher as the time proceeds. Hence, we use a function for weight w over time, giving more importance to the variance criterion as exploration time proceeds:

$$w(t) = \frac{t}{t + r}, \quad (3)$$

where $r \geq 0$ is a constant that tunes the steepness of the curve.

We evaluated both the techniques to generate a set of candidate locations mentioned in Figure 2 using a simulated world with chlorophyll measurements. The details about the simulations will be discussed in Section IV-A. We built a representation of the world as the robot was traveling and collecting measurements. Figures 3 and 4 show the root mean squared error (RMSE) of the phenomenon model

$\hat{f}(\mathbf{x})$ with the ground truth $f(\mathbf{x})$ plotted against the distance traveled. Figure 3 presents a comparison between differently sized fixed windows and Figure 4 presents a comparison between the fixed window and contour-based candidate selection methods. The contour-based method travels longer distances to improve the quality of the map, thus compromising the minimum distance criterion. Because of its better performance, we will consider the fixed window candidate selection technique in the rest of this paper.

B. Look-back Selective Sampling

The explorer robot, while exploring the region to build the model, communicates potential candidates for collecting water samples. Then the sampler's task is to maintain a list of these candidate locations and strategically decide on k locations to collect water samples from – recall that k is the maximum number of samples that can be collected. There are at least two approaches that one can think of for this scenario: the first one is to start making decisions as soon as candidate suggestions come from the explorer; the second is to wait until the explorer has completely mapped the scalar field and then use all the candidate locations to pick k locations. However, in our application of sampling water from the surface of the water body based on its current properties, it is very important that the water quality measurement and the physical water sample are gathered temporally close to each other. This is because of the dynamic behavior of the phenomenon that we are trying to capture. Hence, in this paper we discuss a technique to collect physical samples in parallel with the explorer and achieve a good sampling score by collecting samples within the peaks (*hotspots*) in the spatial field.

Given M measurements – i.e., candidate sampling positions – we need to choose k sample locations that optimize the quality of the final result. Since we are looking at simultaneous decision making along with the explorer, there is a need for optimal stopping criteria – in other words, when the sampler decides to use one of the remaining water sampling units to collect a physical sample. This problem has similarities with the classic *Secretary Problem* that uses optimal stopping theory. The basic form of the secretary problem has n applicants who are interviewed in random order, and a decision is to be made immediately after every interview. Once rejected, an applicant cannot be recalled. So, the problem is to choose an optimal stopping rule to maximize the probability of selecting the best candidate.

Our problem is a variant of this problem as we need to choose k sample points instead of just one. Moreover, we have an advantage: the robot can look back and choose an old candidate if it is the best location to sample water from. In our case, we want to maximize the sum of the values at sampled locations ($\sum_{l \in \mathcal{L}^*} f(l)$) with a minimum distance constraint (T_d). The threshold T_d prevents acquisition of spatially neighboring samples. The value for T_d is application specific and also depends on the possible error in robot localization. We still need a stopping rule to make our decision. Hence, we use a variant of the standard secretary

problem algorithm that suggests we reject the first n/e candidates and then stop at the first candidate with a higher ranking than all the ones evaluated until now. In this way, the probability of success is maximized and is $1/e$ [22]. In our case, we need to choose k samples, hence the stopping threshold becomes $n/(ke)$.

Algorithm 1 Look-back Selective Sampling Algorithm

Input: Number of water sampling units k
Measurements frequency in Hz, f
Mission duration in seconds, T_m
Distance threshold in meters, T_d

Output: List of selected candidates \mathcal{L}^* where sampler should take samples

```

1:  $\tau = \frac{T_m * f}{k}$   $\triangleright$  Total maximum number of measurements
   for each water sampling unit
2:  $\mathcal{L} = \emptyset$   $\triangleright$  List of candidates suggested by the explorer
3:  $\mathcal{L}^* = \emptyset$   $\triangleright$  List of selected candidates
4:  $C_c = 0$   $\triangleright$  Current candidate counter
5: found = false  $\triangleright$  Flag to identify sample chosen within
    $\tau$ 
6: repeat
7:    $l = \text{receiveMeasurementFromExplorer}()$ 
8:    $C_c = C_c + 1$ 
9:   if (Distance( $l, \mathcal{L}^*$ ) >  $T_d$ ) then
10:     $\mathcal{L} = \mathcal{L} \cup l$ 
11:   end if
12:   if  $C_c == \tau/e$  then
13:      $y_{\max} = \max_{l \in \mathcal{L}}(y_l)$   $\triangleright y_l$ : measured value at  $l$ 
14:   else if  $C_c > \tau/e$  then
15:     if  $y_l > y_{\max}$  then
16:        $l^* = l$ 
17:       found = true
18:     else if  $C_c == \tau$  then  $\triangleright$  Time slot expired for  $k$ 
19:        $l^* = \arg \max_{l \in \mathcal{L}} y_l$ 
20:       found = true
21:     end if
22:     if found == true then
23:       goToAndSample( $l^*$ )
24:        $\mathcal{L}^* = \mathcal{L}^* \cup l^*$ 
25:       remove  $l^*$  and its neighbors within  $T_d$  from
          $\mathcal{L}$ 
26:        $C_c = 0$ 
27:       found = false
28:     end if
29:   end if
30: until  $|\mathcal{L}^*| \geq k$ 

```

Kleinberg [24] suggested an algorithm that works by splitting the candidates in approximately two half intervals chosen randomly using a binomial distribution $B(n, 0.5)$. Then, the algorithm proceeds by recursively applying the classic secretary algorithm. The submodular secretary algorithm proposed by Bateni *et al.* [12] provides a mechanism to select the set of candidates with the highest cumulative rating. This algorithm could be a good fit for our problem of




Fig. 5. (a) Gazebo simulation for the ASV [23] used in our experiments. (b) Simulated chlorophyll density map overlaid on top of the water surface. The colorbar shows the simulated chlorophyll density. (c) Part of the Lake Nighthorse, Durango used for running simulation experiments.

making online decision about the water sampling. However, the submodular algorithm splits the samples uniformly into k equal windows and the samples from one window are not considered while making a decision for another window. We would like to have an option to look back into all the windows while making the decision. This is because, in our application, we are not bound by the trajectory followed by the explorer. The sampler robot is free to go back and visit any old measured location if it does not find any eligible candidate as the time proceeds. Hence, we propose a look-back selective sampling technique, where the robot appends new candidates into a list and uses the list to look back if there are no eligible candidates within the time threshold. The pseudo-code for our approach is presented in Algorithm 1.

Line 1 divides the maximum number of measurements over the mission in k uniform time slots. Line 9 checks if a new measurement taken at location l is far enough from the currently selected sampling locations. Line 12 is the secretary problem threshold to get optimal probability. Line 15 makes an irrevocable decision, following the secretary problem algorithm. However, if the time slot (τ) for a water sampling unit expires (Line 18), then the sampler samples from the location with value y_{\max} . The rationale is that given k water sampling units, full mission time (T_m) should not be spent on just one unit.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

We implemented the proposed algorithms in the Robot Operating System (ROS) framework [25] and we evaluated the system both in simulation and in the field on real robots. The simulation environment allows us to report repeatable controlled measurements of performance using realistic data with perfect ground truth. The field deployment lets us observe the performance and feasibility of our approach in practice and confirm its utility and usability. Three different setups are used in this paper to extensively evaluate the proposed system: 1. Simulated robots exploring and sampling from a synthetically created world, 2. Real world data (chlorophyll concentration in the flood plains of the Amazon) is used to create a world for simulated robots, and 3. deployment of two robotic boats in a reservoir to map the




Fig. 6. Chlorophyll concentration dataset used for validating our experiments : (a) Flood plains of Amazon river. Considered region is approximately 2 km \times 2 km (b) Chlorophyll concentration (mg/m^3) map generated from MODIS reflectance values [27] for the area in (a).

chlorophyll density distribution and collect water samples rich in chlorophyll content.

A. Simulations

Gazebo with an ASV plugin [26] that simulates a physically realistic Clearpath Heron robotic boat (Figure 5(a)) is used in our experiments due to its capabilities of simulating the vehicle dynamics to reasonable precision. We created a ROS node to simulate a water quality sensor, returning measurements at given GPS locations, according to some data source.

The synthetic world, along with the simulator, are shown in Figure 5. In particular, Figure 5(b) shows synthetic data simulating a chlorophyll density field overlaid on top of the water surface. To generate such data, we used multi-Gaussian models to imitate the chlorophyll dense regions and its diffusion on the water surface. GPS data from a region (100 m \times 90 m) in Lake Nighthorse, Durango, CO (shown in Figure 5(c)) is used as the underlying localization for our simulations. We performed five repetitions for each of the experiments in real time simulation.

The real setup in our simulations uses the chlorophyll concentration ($\mu\text{g/m}^3$) map, at the flood plains of Amazon, generated from MODIS reflectance values [27]. We chose a bounded region of size 2 km \times 2 km (Figure 6(a)) from this dataset to build an environment for the simulated ASVs. The spatial field from this dataset is presented in Figure 6(b). We evaluate our system by initially comparing and testing both the components (*explorer and sampler*) separately and then we present results from the whole system coordinating




Fig. 7. Trajectories planned and traversed (Yellow lines) by the ASV according to: (a) Planning window based GP-frontier explorer. (b) Global maximum variance search. (c) Lawnmower coverage.

towards the collection of water samples with high utility measured in terms of sampling score.

1) *Evaluating the explorer*: To estimate the utility of our exploration algorithm, we measure its performance in terms of the root mean squared error (RMSE) between the generated representation and the ground truth over distance traveled and compare this to alternative state of the art approaches. Note that most methods assure good data as time (or distance traveled) approaches infinity, but one attribute of interest is to try and acquire a good estimate as early as possible. Prior approaches to such coverage and sampling problems can be grouped as deterministic complete coverage (such as the boustrophedonic coverage, or “lawnmower” algorithm [4] or stochastic methods).

In this paper, we compare the GP-frontier based explorer to two other exploration techniques: global maximum variance search, and lawnmower coverage. Global maximum variance search involves predicting the variance at every location in the region and then searching over the entire grid world. These two operations are computationally expensive compared to a small set of predictions needed for our approach. As a reminder, our approach needs predictions only at the locations that lie on the planning window boundary (Figure 2(b)). Furthermore, global maximum variance search generates longer trajectories as shown in Figure 7(b) thus making it power inefficient compared to our approach. The traditional approach to covering a partially observable, obstacle-free region is to employ a *boustrophedonic* or *lawnmower coverage*. Such complete surveys (Figure 7(c)) are infeasible as we are limited by battery life on the robots.

We compare these exploration techniques by computing the RMSE of the generated representation relative to the ground truth data (Figure 5(b) and Figure 6(b)):

$$\text{RMSE}_d = \sqrt{\frac{\sum_{c \in \mathcal{E}'} (\hat{f}^t(c) - f(c))^2}{|\mathcal{E}'|}}, \quad (4)$$

where RMSE_d is the root mean squared error of the generated map of Chlorophyll density after traveling a given distance d , $c \in \mathcal{E}'$ is the set of cells from the discretized world, $\hat{f}^t(c)$ is the predicted value at c with the GP at time t , and $f(c)$ is the ground truth value at c . The plots in Figure 8 illustrate that, as the travel distance increases, more of the world is explored by all techniques and they all

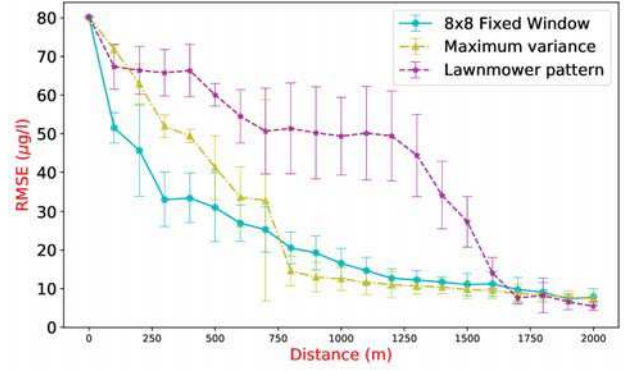


Fig. 8. RMSE of the generated map of Chlorophyll density plotted against travel distance. Plots from different exploration techniques validated on simulated data.




Fig. 9. RMSE of the generated map of Chlorophyll density plotted against travel distance. Plots from different exploration techniques validated on real Chlorophyll concentration data from Amazon flood plains.

converge to a good representation of the world. Nevertheless, the GP-frontier explorer generates a good representation of the spatial field with less traveled distance, by choosing right locations to visit and map, thus providing better results than other techniques compared in this paper.

Figure 9 presents the RMSE when operating in a larger field and operating in an environment created using the real chlorophyll concentration dataset. The GP-frontier based explorer performs well in the beginning and later performs on par with the maximum variance search technique. The larger field and fast changing weight function (Eq. (2)) affect the performance of our technique.

2) *Evaluating the sampler*: A smart sampler can choose locations with a spectrum of measurements to represent the diversity in the spatial field, or it can sample from




Fig. 10. Simulated distributions of chlorophyll density with the candidate locations (red dots) selected for collecting water samples. (a) Look-back Selective Sampling approach, (b) Submodular Secretary Algorithm.




Fig. 11. (a) Rogers Reservoir, Durango, CO, where field experiments were conducted. (b) The explorer trajectory (yellow) and the chosen sampling candidates (red circles). (c) The spatial mapping of chlorophyll density ($\mu\text{g/l}$) generated using explorer measurements. Red-dots are the chosen candidate locations for water samples.

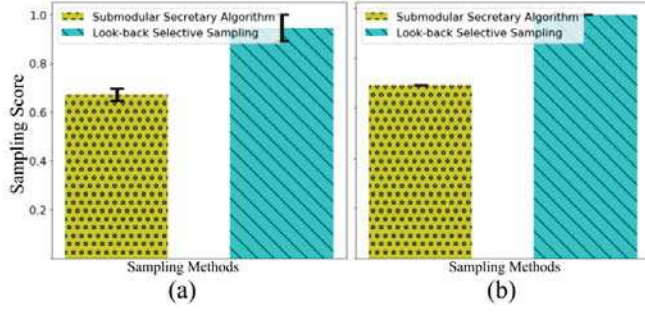


Fig. 12. Results from simulations with synthetic data (a) and real data (b). The bar plots indicate the sampling score between two sampling methods. The error bars show the standard deviation over five real time trials.

locations that give high rewards. For our application, we want the samples from *hotspot* regions that are high in chlorophyll concentration. However, we also do not want all the samples to be from the same spot. Hence, to evaluate such a system we propose a sample scoring metric which evaluates the sampling techniques according to their ability to choose non-neighboring samples from hotspot regions. The maximum value (M_{value}) that can be achieved by any sampling technique is computed by summing the k largest values among all data measurements ($f(\mathbf{x})$) provided by the explorer. The scoring function is the ratio of value achieved by the sampling algorithm to the maximum achievable value:

$$\text{Score} = \frac{\sum_{l \in \mathcal{L}^*} f(l)}{M_{\text{value}}}. \quad (5)$$

We compare our look-back selective sampling with the submodular secretary algorithm [12] and Figure 10 presents the sample locations chosen (red dots) by both the algorithms. Looking at the score of the two algorithms (Figure 12), the look-back selective sampling technique appears to be more suitable for our application. This is because the submodular secretary algorithm divides the whole segment into windows and considers the candidates only within a window. However, in our application of sampling from a bounded region, we are not constrained by not being able to go back spatially to take a sample.

Figure 13 illustrates the performance of the whole system, explorer and the sampler working together to achieve good sample quality. We conducted a series of experiments with three explorer-sampler pairs. The results show that the mul-




Fig. 13. Sampling scores achieved by the complete system, using different combinations of explorers and the look-back selective sampler on synthetic data.




Fig. 14. Results from field experiments. Comparison of sampling scores between the two sampling methods.

tirobot system with the proposed components performs well by achieving samples with high sampling scores.

B. Field experiments

Our goal is to sample water from a closed water body for ex-situ analysis. We are interested in getting samples that are rich in chlorophyll. Rogers Reservoir, located in Durango, Colorado (Figure 11(a)), is a drinking-water reservoir and monitoring this water body is very essential for the city. For our field experiments, we used two Clearpath Heron vehicles (see Figure 1): one equipped with a water quality sensor that collects measurements at 1 Hz – the explorer – and another equipped with a water sampling apparatus – the sampler. Figures 11(b) and (c) show the trajectory followed by the explorer and a reconstruction of chlorophyll distribution in the region of survey. The candidate locations chosen by the sampler are shown with red circles. The field experiments also confirmed that the proposed look-back sampling technique achieves higher sampling scores (Figure 14).

It is worth to mention some of the additional issues that

need to be considered during field experiments: GPS errors; the fact that a dense coverage should be run every time before the actual experiment to ensure that we have the most recent data as ground truth to evaluate our method. Especially in the marine domain, weather can affect the schedule and also the properties of the environment.

V. CONCLUSIONS

In this paper, we proposed a heterogeneous multirobot system for physical sampling of a water body providing methods for two related subproblems: one exploration algorithm to build the phenomenon map, which concurrently drives the sampling algorithm to actually collect physical samples. The core of the approach is to combine efficient real time estimation of a variable upon which our phenomenon of interest is conditionally dependent with the subsequent collection of data. Our approach allows us to efficiently collect a set of informative samples lowering the uncertainty over it and sample from significant hotspots, as compared to other methods. This is validated through simulations and the feasibility and practicality of the system is demonstrated via field experiments.

With respect to future and ongoing work, we are scaling up the approach for application over larger regions in more challenging outdoor environments. This entails the use of faster and more capable marine vehicles. Explicitly modeling and accounting for communication interruption in the decision making process is also an important next step to ensure reliability and robustness over large regions of space and time [28], [29]. The consideration of time-varying models will also be an interesting step towards more large-scale deployment in marine environments.

REFERENCES

- [1] J. Das, T. Maughan, M. McCann, M. Godin, T. O'Reilly, M. Messi, F. Bahr, K. Gomes, F. Py, J. G. Bellingham, G. S. Sukhatme, and K. Rajan, "Towards mixed-initiative, multi-robot field experiments: Design, deployment, and lessons learned," in *Proc. IROS*, 2011, pp. 3132–3139.
- [2] J. Das, F. Py, J. B. J. Harvey, J. P. Ryan, A. Gellene, R. Graham, D. A. Caron, K. Rajan, and G. S. Sukhatme, "Data-driven robotic sampling for marine ecosystem monitoring," *Int. J. Robot. Res.*, vol. 34, no. 12, pp. 1435–1452, 2015.
- [3] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proc. International Conference on Autonomous Agents*, 1998, pp. 47–53.
- [4] H. Choset and P. Pignon, *Coverage Path Planning: The Boustrophedon Cellular Decomposition*. London: Springer London, 1998, pp. 203–209.
- [5] A. Xu, C. Viriyasuthee, and I. Rekleitis, "Efficient complete coverage of a known arbitrary environment with applications to aerial operations," *Auton. Robot.*, vol. 36, no. 4, pp. 365–381, 2014.
- [6] S. Manjanna, N. Kakodkar, M. Meghiani, and G. Dudek, "Efficient terrain driven coral coverage using Gaussian processes for mosaic synthesis," in *Proc. CRV*, 2016, pp. 448–455.
- [7] S. Manjanna and G. Dudek, "Data-driven selective sampling for marine vehicles using multi-scale paths," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vancouver, Canada, September 2017, pp. 6111–6117.
- [8] K. H. Low, J. M. Dolan, and P. Khosla, "Adaptive multi-robot wide-area exploration and mapping," in *Proc. AAMAS*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 23–30.
- [9] G. Hitz, A. Gotovos, M.-É. Garneau, C. Pradalier, A. Krause, R. Y. Siegwart, et al., "Fully autonomous focused exploration for robotic environmental monitoring," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 2658–2664.
- [10] M. Rahimi, M. Hansen, W. J. Kaiser, G. S. Sukhatme, and D. Estrin, "Adaptive sampling for environmental field estimation using robotic sensors," in *Proc. IROS*, 2005, pp. 3692–3698.
- [11] Y. Girdhar and G. Dudek, "Optimal online data sampling or how to hire the best secretaries," in *Proc. CRV*, 2009, pp. 292–298.
- [12] M. Bateni, M. Hajiaghayi, e. M. Zadimoghaddam, Morteza, R. Shaltiel, K. Jansen, and J. Rolim, "Submodular secretary problem and extensions," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010*. Springer Berlin Heidelberg, 2010, pp. 39–52.
- [13] A. Valada, P. Velagapudi, B. Kannan, C. Tomaszewski, G. Kantor, and P. Scerri, "Development of a low cost multi-robot autonomous marine surface platform," in *Field and Service Robotics: Results of the 8th International Conference*, K. Yoshida and S. Tadokoro, Eds. Springer Berlin Heidelberg, 2014, pp. 643–658.
- [14] Y. Girdhar, A. Xu, B. B. Dey, M. Meghiani, F. Shkurti, I. Rekleitis, and G. Dudek, "MARE: Marine Autonomous Robotic Explorer," in *Proc. IROS*, 2011, pp. 5048–5053.
- [15] M. Caccia, R. Bono, G. Bruzzone, E. Spirandelli, G. Veruggio, A. M. Stortini, and G. Capodaglio, "Sampling sea surfaces with SESAMO: an autonomous craft for the study of sea-air interactions," *IEEE Robot. Automat. Mag.*, vol. 12, no. 3, pp. 95–105, 2005.
- [16] J.-P. Ore, S. Elbaum, A. Burgin, and C. Detweiler, "Autonomous aerial water sampling," *J. Field Robot.*, vol. 32, no. 8, pp. 1095–1113, 2015.
- [17] P. S. Schenker, T. L. Huntsberger, P. Pirjanian, E. T. Baumgartner, and E. Tunstel, "Planetary Rover Developments Supporting Mars Exploration, Sample Return and Future Human-Robotic Colonization," *Auton. Robot.*, vol. 14, no. 2, pp. 103–126, 2003.
- [18] N. Basilico and F. Amigoni, "Exploration strategies based on multi-criteria decision making for searching environments in rescue operations," *Auton. Robot.*, vol. 31, no. 4, pp. 401–417, 2011.
- [19] A. Lilienthal, A. Zell, M. Wandel, and U. Weimar, "Sensing odour sources in indoor environments without a constant airflow by a mobile robot," in *Proc. ICRA*, vol. 4, 2001, pp. 4005–4010.
- [20] A. Quattrini Li, R. Cipolleschi, M. Giusto, and F. Amigoni, "A semantically-informed multirobot system for exploration of relevant areas in search and rescue settings," *Auton. Robot.*, vol. 40, no. 4, pp. 581–597, 2016.
- [21] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [22] E. B. Dynkin, "The optimum choice of the instant for stopping a markov process," in *Soviet Math. Dokl*, vol. 4, no. 627–629, 1963.
- [23] N. Koenig and A. Howard, "Design and use paradigms for Gazebo, an open-source multi-robot simulator," in *Proc. IROS*, vol. 3, pp. 2149–2154.
- [24] R. Kleinberg, "A multiple-choice secretary algorithm with applications to online auctions," in *Proc. ACM-SIAM symposium on Discrete algorithms*, 2005, pp. 630–631.
- [25] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, 2009.
- [26] B. Bingham, "Unmanned Surface Vehicle plugins for Gazebo simulation," https://github.com/bsb808/usv_gazebo_plugins, 2017, commit: 0a01cab47ebf74aad75775a5d4eaf727be7bfc77.
- [27] E. M. L. de Moraes Novo, C. C. de Farias Barbosa, R. M. de Freitas, Y. E. Shimabukuro, J. M. Melack, and W. Pereira Filho, "Seasonal changes in chlorophyll distributions in amazon floodplain lakes derived from modis images," *Limnology*, vol. 7, no. 3, pp. 153–161, 2006.
- [28] J. Banfi and A. Quattrini Li and N. Basilico and I. Rekleitis and F. Amigoni, "Multirobot online construction of communication maps," in *Proc. ICRA*, 2017, pp. 2577–2583.
- [29] J. Banfi, A. Quattrini Li, N. Basilico, I. Rekleitis, and F. Amigoni, "Asynchronous multirobot exploration under recurrent connectivity constraints," in *Proc. ICRA*, 2016, pp. 5491–5498.